# PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named' in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed

Dated   30 October 2003

Patents Form 1/77

Patents Act 1977
(Rule 16)

THE PATENT OFFICE

19 JUL 2002

The Patent Office

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

22JUL02 E734821-2 002847
P01/7700 0.00-0216880.5

The Patent Office

Cardiff Road
Newport
South Wales
NP10 8QQ

1. Your reference	HL82518/000/DCO	
2. Patent application number *(The Patent Office will fill in this part)*	19 JUL 2002	**0216880.5**
3. Full name, address and postcode of the or of each applicant *(underline all surnames)*	PICOCHIP DESIGNS LIMITED Second Floor Suite, Riverside Buildings 108 Walcot Street Bath BA1 5BG	

08047151002

Patents ADP number *(if you know it)*

If the applicant is a corporate body, give the country/state of its incorporation

United Kingdom

4. Title of the invention

PROCESSOR ARRAY

5. Name of your agent *(if you have one)*

Haseltine Lake

"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*

Imperial House
15-19 Kingsway
London
WC2B 6UD

Patents ADP number *(if you know it)*   34001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and *(if you know it)* the or each application number

Country	Priority application number *(if you know it)*	Date of filing *(day / month / year)*

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application	Date of filing *(day / month / year)*

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? *(Answer 'Yes' if*
   a) *any applicant named in part 3 is not an inventor, or*
   b) *there is an inventor who is not named as an applicant, or*
   c) *any named applicant is a corporate body.*
   See note (d))

YES

Patents Form 1/77

Patents Form 1/77

9. Enter the number of sheets for any of the
   following items you are filing with this form.
   Do not count copies of the same document

Continuation sheets of this form	0
Description	17
Claim(s)	3 C.F
Abstract	1
Drawing(s)	15 + 15

10. If you are also filing any of the following,
    state how many against each item.

Priority documents	-
Translations of priority documents	-
Statement of inventorship and right to grant of a patent *(Patents Form 7/77)*	-
Request for preliminary examination and search *(Patents Form 9/77)*	1
Request for substantive examination *(Patents Form 10/77)*	-
Any other documents *(please specify)*	-

11.   I/We request the grant of a patent on the basis of this application.

Signature   Date

Haelthe Lahle

12. Name and daytime telephone number of person to contact in the United Kingdom   Mrs J. McCormick   [0117] 910 3200

Warning
After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes
a) *If you need help to fill in this form or you have any questions, please contact the Patent Office on 08459 500505.*

b) *Write your answers in capital letters using black ink or you may type them.*

c) *If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*

d) *If you have answered 'Yes' Patents Form 7/77 will need to be filed.*

e) *Once you have filled in the form you must remember to sign and date it.*

f) *For details of the fee and ways to pay please contact the Patent Office.*

Patents Form 1/77

PROCESSOR ARRAY

This invention relates to a processor array, and in particular to a processor array with a degree of

5   redundancy which allows the array to operate normally, even in the presence of one or more defective processor.

GB-A-2370380 discloses a processor array, in which data
10   processing functions are distributed amongst processors in an array, the processors being linked by buses and switch elements which determine how data is transferred from one array element to another.

15   Manufacturing processes for semiconductor devices are imperfect. These imperfections result in point defects that are distributed over a silicon wafer. For a given defect density, if the die size is larger, then the proportion of devices with defects will be greater.
20   For most semiconductor devices, if a defect occurs anywhere on the die then that die must be discarded, because all the circuitry in the device is required for correct operation.

25   One known exception to this is in the case of memory devices, such as Random Access Memories (RAMs). In this case, because the bulk of the device consists of a regular array of memory cells, spare (redundant) columns of cells may be incorporated in the device that
30   can be used to replace columns in which defects are detected during testing. In order to achieve the replacement of defective columns, switches, controlled by means of laser fuses, are incorporated in the circuitry. These fuses are selectively blown as a

2

result of information obtained from the testing. This
increases the proportion of usable devices that can be
obtained from a wafer.

5   The present invention provides an array of processing
elements, which can incorporate a degree. of redundancy.
Specifically, the array includes one or more spare, or
redundant, rows of array elements, in addition to the
number required to implement the intended function or
10   functions of the device. If a defect occurs in one of
the processors in the device, then the entire row which
includes that defective processor is not used, and is
replaced by a spare row.

15   According to a first aspect of the present invention,
- there is provided method of replacing a faulty
processor element, in a processor array comprising a
plurality of processor elements arranged in an array of
rows and columns, the processor elements being
20   interconnected by buses running between the rows and
columns and by switches located at the intersections of
the buses, and the array including a redundant row to
which no functionality is initially allocated. In the
event that a first processor element is found to be
25   faulty, functionality is removed from the row that
contains said first processing element, and allocated
instead to the redundant row.

This allows the required functionality to be carried
30   out, even on a device which includes a faulty processor
element. This can significantly increase the
proportion of usable devices which are obtained from
the manufacturing process.

According to a second aspect of the present invention, there is provided a processor array, which has processor elements arranged in an array of rows and columns, wherein the arrangement of processor elements

5 in each row is the same as the arrangements of processor elements in each other row; pairs of horizontal buses running between the rows of processor elements, each pair comprising a first horizontal bus carrying data in a first direction and a second

10 horizontal bus carrying data in a second direction opposite to the first direction; vertical buses running between the columns of processor elements, wherein some pairs of adjacent columns of processor elements have no vertical buses running therebetween, and other pairs of

15 adjacent columns have two buses carrying data in a first direction and two buses carrying data in a second direction opposite to the first direction running therebetween; and switches located at the intersections of the horizontal and vertical buses.

20

This array, and in particular the uneven arrangement of the vertical buses, allows the most efficient use of the method according to the first aspect of the invention.

25

For a better understanding of the present invention, and to show how it may be put into effect, reference will now be made, by way of example, to the accompanying drawings, in which:-

30

Figure 1 is a block schematic diagram of an array according to the present invention.

Figure 2 is a block schematic diagram of a switch within the array of Figure 1.

Figure 3 is an enlarged block schematic diagram of a
5   part of the array of Figure 1.

Figure 4 is a schematic representation of a semiconductor wafer used to manufacture the array of Figure 1.

10

Figure 5 is a block schematic diagram of the array of Figure 1, showing the effect of a possible defect.

Figures 6-16 are enlarged block schematic diagrams of a
15   part of the array of Figure 1, showing the operation of the device in the event of possible defects.

Figure 1 shows an array architecture in accordance with the present invention. The array architecture is
20   generally as described in GB-A-2370380 and GB-A-2370381, which are incorporated herein by reference, with modifications which will be described further herein.

25   The array consists of a plurality of array elements 20, arranged in a matrix. For ease of illustration, the example shown in Figure 1 has six rows, each consisting of ten array elements (AE0, AE1, ..., AE9), giving a total of 60 array elements, but a practical embodiment
30   of the invention may for example have over 400 array elements in total. Each array element, 20, is connected to a segment of a respective first horizontal bus 32 running from left to right, and to a segment of a respective second horizontal bus 36 running from

right to left, by means of respective connectors 50.
The horizontal bus segments 32, 36 are connected to
vertical bus segments 41, 43 running upwards and to
vertical bus segments 42, 44 running downwards, at
5   switches 55. Specifically, each switch 55 has an input
left-right horizontal bus segment 32, an input right-
left horizontal bus segment 36, two input upwards
vertical bus segments 41, 43, and two input downwards
vertical bus segments 42, 44, plus an output left-right
10   horizontal bus segment 32, an output right-left
horizontal bus segment 36, two output upwards vertical
bus segments 41, 43, and two output downwards vertical
bus segments 42, 44.

15   All horizontal bus segments 32, 36 and vertical bus
segments 41, 43, 42, 44 are 32 bits wide.

Thus, while some pairs of adjacent columns of processor
elements (e.g. AE1 and AE2, AE6 and AE7) have no
20   vertical buses running therebetween, other pairs of
adjacent columns (e.g. AE4 and AE5, AE8 and AE9) have
two buses carrying data in a first direction and two
buses carrying data in a second direction opposite to
the first direction running therebetween. This
25   unevenly spaced arrangement, providing two pairs of
vertical buses after a group of four columns of array
elements rather than, say, one pair of vertical buses
after a group of two columns of array elements, is more
efficient, for reasons which will be described below.
30

Figure 2 shows the structure of one of the switches 55,
each of the switches being the same. The switch
includes a random access memory RAM 61, which is pre-
loaded with data. The switch 55 is locally controlled

by a controller 60, which contains a counter that
counts through the addresses of RAM 61 in a pre-
determined sequence. This same sequence is repeated
indefinitely, and the time taken to complete the

5   sequence once, measured in cycles of the system clock,
is referred to as the sequence period. On each clock
cycle, the output data from RAM 61 is loaded into a
register 62, and the content of the register 62 is used
to select the source for each output bus 66 using

10 multiplexer 64.

The source for the output bus 66 may be any one of the
six input buses, namely the input left-right horizontal
bus segment LeftIn, the input right-left horizontal bus

15 segment RightIn, the two input upwards vertical bus
segments Up1In and Up2In, or the two input downwards
vertical bus segments Down1In and Down2In. In
addition, the value zero can be selected as the source
for an output bus, as can the value that was on the

20 output bus during a previous clock cycle, which is
loaded into a register 65 under the control of one of
the bits in register 62.

When an output bus is not being used, it is

25 advantageous to select zero as the source, so that the
value on the bus will remain unchanged over several
clock cycles, thereby conserving power.

In Figure 2, only one multiplexer 64 and its associated

30 register 65 is shown, although the switch 55 includes
six such multiplexers and associated registers, one for
each output bus.

The biggest component of the switch 55 is the RAM 61,
although this is still small by the standards of RAMs
generally. Therefore, the size of the RAM 61 is
dictated to a large extent by the address decoding

5   section of the RAM. Since this is not dependent on the
number of buses being switched in the switch 55, the
overall size of the device can be reduced by providing
two pairs of vertical buses, and one switch in each
row, after a group of four columns of array elements,

10   as compared with providing, say, one pair of vertical
buses, and one switch in each row, after each group of
two columns of array elements.

Figure 3 shows in more detail how each array element 20

15   is connected to the segments of the horizontal buses
32, 36 at a connector 50. Each such connector 50
includes a multiplexer 51, and a segment of the bus is
defined as the portion between two such multiplexers
51. Each segment of a left-right horizontal bus 32 is

20   connected to an input of a respective array element 20
by means of a connection 21. An output 22 of each
array element 20 is connected to a respective segment
of a right-left horizontal bus 36 through another
multiplexer 51.

25

Each multiplexer 51 is under control of circuitry (not
shown) within the associated array element 20, which
determines whether the multiplexer outputs the data on
the input bus segment or the data on the array element

30   output.

All communication within the array takes place in a
predetermined sequence, which lasts for a predetermined
number (for example, 1024) of clock cycles (the

sequence period that is described above). Each switch
and each array element contains a counter that counts
for the sequence period. As described above, on each
cycle of this sequence, each switch selects data from
5   one of the eight possible sources onto each of its six
output buses. At predetermined cycles, array elements
load data in from the respective input bus segments via
the respective connections 21, and switch data onto the
respective output bus segments, using the multiplexers
10   51.

Each array element is capable of controlling its
associated multiplexer, and loading data from the bus
segments to which it is connected at the correct times
15   in sequence, and of performing some useful function on
the data. The useful function may consist only of
storing the data.

However, in a preferred embodiment of the invention,
20   each array element contains a complex microprocessor,
the area of which is several times more than that of
each switch. This difference in size makes the present
invention, which is concerned with overcoming failures
in the array elements rather than in the switches,
25   particularly effective.

If the array elements are not all identical then, in
order for the present invention to be usable most
efficiently, at least within each column of the array,
30   all of the array elements should be identical.

The manufacturing processes for semiconductor devices
are imperfect. This results in point defects that are
distributed over a silicon wafer. For a given

manufacturing process at a given state of maturity this defect density will be roughly constant. This is illustrated in Figure 4, which shows the boundary of a circular silicon wafer 60, the individual square dice

5   61 which are used for making individual devices, and randomly distributed defects 62.

For a given defect density, if the die size is larger, then the proportion of devices with defects will be

10 greater.

For most semiconductor devices, if a defect occurs anywhere on the die then that die must be discarded, because all the circuitry in the device is required for

15 correct operation.

According to the present invention, the array of processing elements incorporates a degree of redundancy. More specifically, one or more spare

20 (redundant) rows of array elements, over and above the number required to implement the intended function or functions of the device, is included in the array. If a defect occurs in one of the processing elements, either during manufacturing or in operation of the device,

25 then the entire row of array elements which includes the defective processing element is not used, and is replaced by a spare row.

Since the array elements within a column are all

30 identical, as mentioned above, each row of array elements is identical, and all of the functionality of the row that includes the defective processing element can be performed by the spare row.

Figure 5 shows an array of processing elements as shown in Figure 1, with the six rows indicated as Row n (n = 0, ..., 5), and the ten array elements in each row indicated as AEnm (m= 0, ..., 9). If a defect is

5 detected in any of the array elements in row 2, for example, then that entire row of array elements is not used. In more detail, if it was originally intended that the spare row should be Row 5, and the production test process detects a fault in an array element in Row

10 2, then the software programs that would otherwise have been loaded into the array elements of Row 2 are loaded into the corresponding array elements of Row 3; the software programs that would otherwise have been loaded into the array elements of Row 3 are loaded into the

15 corresponding array elements of Row 4; and the software programs that would otherwise have been loaded into the array elements of Row 4 are loaded into the corresponding array elements of Row 5.

20 Denoting the software program that was originally destined for any one of the array elements AEnm as Prognm, the complete redistribution of programs is defined below:

Array Element	Original Program	Revised Program	Array Element	Original Program	Revised Program	Array Element	Original Program	Revised Program
AE00	Prog00	Prog00	AE01	Prog01	Prog01	AE02	Prog02	Prog02
AE03	Prog03	Prog03	AE04	Prog04	Prog04	AE05	Prog05	Prog05
AE06	Prog06	Prog06	AE07	Prog07	Prog07	AE08	Prog08	Prog08
AE09	Prog09	Prog09	AE10	Prog10	Prog10	AE11	Prog11	Prog11
AE12	Prog12	Prog12	AE13	Prog13	Prog13	AE14	Prog14	Prog14
AE15	Prog15	Prog15	AE16	Prog16	Prog16	AE17	Prog17	Prog17
AE18	Prog18	Prog18	AE19	Prog19	Prog19	AE20	Prog20	Unused
AE21	Prog21	Unused	AE22	Prog22	Unused	AE23	Prog23	Unused

AE24	Prog24	Unused	AE25	Prog25	Unused	AE26	Prog26	Unused
AE27	Prog27	Unused	AE28	Prog28	Unused	AE29	Prog29	Unused
AE30	Prog30	Prog20	AE31	Prog31	Prog21	AE32	Prog32	Prog22
AE33	Prog33	Prog23	AE34	Prog34	Prog24	AE35	Prog35	Prog25
AE36	Prog36	Prog26	AE37	Prog37	Prog27	AE38	Prog38	Prog28
AE39	Prog39	Prog29	AE40	Prog40	Prog30	AE41	Prog41	Prog31
AE42	Prog42	Prog32	AE43	Prog43	Prog33	AE44	Prog44	Prog34
AE45	Prog45	Prog35	AE46	Prog46	Prog36	AE47	Prog47	Prog37
AE48	Prog48	Prog38	AE49	Prog49	Prog39	AE50	Unused	Prog40
AE51	Unused	Prog41	AE52	Unused	Prog42	AE53	Unused	Prog43
AE54	Unused	Prog44	AE55	Unused	Prog45	AE56	Unused	Prog46
AE57	Unused	Prog47	AE58	Unused	Prog48	AE59	Unused	Prog49

In addition to redistributing the programs that run on
the array elements, the contents of the RAMs in the
switches must also be changed, so that the data is
5   transferred to the array element which will now be
using it. As can be seen from the table above, the
programs are redistributed in such a way that programs
which are run on array elements above the row with the
defective array element are run on the same array
10   element, while programs which are run on array elements
in or below the row with the defective array element
are run on the corresponding array element in the row
below their original row. In the same way, therefore,
the RAMs are reprogrammed so that the routes taken by
15   data move down with the failed row and the rows below
it, and stay in the same places in the row above the
failed row. When the routes begin or end on the failed
row or pass through the failed row, the situation is
slightly more complex. All cases will be described
20   with the aid of Figure 6, which shows a detailed view
of part of the array of Figure 5, and also Figures 7 to

16, which show specific examples of the requirements
for switch reprogramming.

As illustrated in Figure 5, it is assumed that one of
5 the array elements in Row2 has failed, and hence that
all of the array elements in that row are no longer to
be used. Specifically, referring to Figure 6, the
programs that would have run on AE24, AE25 and AE26,
are transferred to AE34, AE35, and AE36, respectively.
10 Similarly, the programs that would have run on AE34,
AE35 and AE36, are transferred to AE44, AE45, and AE46,
and so on. In Figures 7 to 16, bold dotted lines
indicate resources (connectors, switches or
multiplexers) that are used in the original (fault-
15 free) configuration, double lines indicate resources
that are used in the new configuration and bold solid
lines indicate resources that are used in both the
original and new configurations.

20 Data is rerouted according to the rules set out below.

For a horizontal route above the failed row then there
is no change.

25 For a horizontal route on the failed row or below it,
all routing is moved down one row. This is illustrated
in Figure 7, which shows that a route from AE24 to
AE25, via bus segments 80 and 81 and switch SW21, is
moved to a route from AE34 to AE35, via bus segments 88
30 and 89 and switch SW31.

Routes lying entirely above the failed row are
unaffected, as illustrated in Figure 8.

Routes lying entirely below the failed row are all moved down by one row, in the same way as horizontal routes, as described above.

5   Routes going to the failed row from below are handled in the same way as routes lying entirely below the failed row and this is illustrated in Figure 9. This shows that a route from AE35 to AE25, via bus segment 95, switch SW41, bus segment 91, switch SW31, bus

10   segment 83, switch SW21, and bus segment 81 is moved to a route from AE45, to AE35, via bus segment 105, switch SW51, bus segment 101, switch SW41, bus segment 91, switch SW31, and bus segment 89.

15   Routes going to the failed row from above, where the original route contains at least one vertical bus segment require that the vertical route must be extended down by one more switch. This is illustrated in Figure 10. This shows a partial route, from switch

20   SW11, to AE25, via bus segment 73, switch SW21, and bus segment 81, which is extended from switch SW21, to AE35, via bus segment 85, switch SW31, and bus segment 89.

25   Routes to the failed row from the row above the failed row that do not use any vertical bus segments form a special case. Specifically, as illustrated in Figure 11, the original route contains no existing vertical bus segments to extend, unlike the example shown in

30   Figure 10. As a result, a new vertical section must be inserted. This could lead to a potential problem, as the vertical bus segment that is required for the new route may have already been allocated to another route during the required time slot.

Therefore, when allocating routes, this requirement must be taken into account. Specifically, when planning any route that goes from one row to another
5 without using any vertical bus segments, the system should reserve a connection on the bus segment that would be required in the event of a failure in the lower of the two rows.

10 The unevenly spaced arrangement of vertical buses, described above, in which two pairs of vertical buses are provided after each group of four columns of array elements, allows these connections to be reserved more efficiently compared with an alternative more even
15 spacing in which, say, one pair of vertical buses is provided after each group of two columns of array elements.

In Figure 11, the route from array element AE15 to
20 array element AE25, via bus segment 75, switch SW21, and bus segment 81, is moved so that bus segment 81 is not used and the route from switch SW21 is extended to array element AE35, via vertical bus segment 85, switch SW31, and bus segment 89. A connection on bus segment
25 85 (or on the other bus segment 82 that runs from switch SW21 to SW31) must be reserved for that time slot by the original route allocation process.

Routes starting from the failed row and going down, as
30 shown in Figure 12, are handled in the same way as routes lying entirely below the failed row. For example, a route from array element AE25 to AE35, via bus segment 87, switch SW31, and bus segment 89, is

moved to a route from array element AE35 to AE45, via bus segment 95, switch SW41, and bus segment 107.

It should be noted that, in this case, although the original route did not include any section of a vertical bus, neither does the replacement route. As a result, this does not lead to the potential problem described above with reference to Figure 11.

Routes coming from the failed row and going up, where the original route includes at least one vertical bus segment, are extended by one vertical bus segment as illustrated in Figure 13. Here, a route from array element AE25 to AE15, via bus segment 87, switch SW31, bus segment 83, switch SW21, bus segment 71 and switch SW11, is altered so that it starts at array element AE35, and is routed to SW31 via bus segment 95, switch SE41, and bus segment 91, thereafter continuing as before to AE15.

A route from the failed row, to the row above, not using any vertical bus segments, is illustrated in Figure 14. This is analogous to the case shown in Figure 11, above. Specifically, a route from array element AE24 to AE14 via bus segment 80, switch SW21 and bus segment 74 is replaced by a route from array element AE34 to AE14 via bus segment 88, switch SW31, bus segment 84, switch SW21 and bus segment 74. Thus, a connection on bus segment 84 (or on the other bus segment 83 that runs from switch SW31 to SW21) must be reserved for that time slot by the original route allocation process, to avoid the possibility of contention.

Routes that cross the failed row are shown in Figure 15
(running upwards) and Figure 16 (running downwards).
In these cases, the allocations of all the vertical bus
segments below the failed row are moved down by one
5   row, and an extra bus segment is allocated in the
failed row.

The process of determining the required contents of the
RAM 61 of each switch, in each of the above cases, will
10   be clear to one skilled in the art.

The above description shows how a single failed row may
be replaced by a spare row. If two spare rows are
included in the array, then two failed rows can be
15   replaced. The process is exactly the same as that
described above, but is repeated twice. First, the
highest failed row is replaced, with the fact that
there is a second failed row being ignored. Then, the
lower failed row is replaced. In principle, any number
20   of failed rows may be repaired in the same way,
although a practical restriction on the number of
replacements is that, since vertical routes may become
longer after each row is replaced (because the routing
is effectively "stretched" over the failed row), this
25   increases the transit time for the data. Eventually,
the increased transit time would mean that the data
could not be processed at the required rate.

If failures are detected as part of production test,
30   the information about which rows contain failures may
be used to blow laser fuses on the devices under test.
The process that loads programs onto the array elements
and manipulates the data in the RAMs of the switches
may use this information at the time the array is

configured.   Alternatively, the method described here
may be used to repair failures that occur during
operation in the field.   In this case, failures in
array elements may be detected by running test software
5   on the array elements.

The method and apparatus are described herein primarily
with reference to an arrangement in which the processor
elements include microprocessors.   As noted above, the
10   processor elements may simply be able to store data.
Conversely, each processor element may itself contain
an array of smaller processor elements.

CLAIMS

1. A method of replacing a faulty processor element, in a processor array comprising a plurality of processor elements arranged in an array of rows and columns, the processor elements being interconnected by buses running between the rows and columns and by switches located at the intersections of the buses, and the array including a redundant row to which no functionality is initially allocated, the method comprising:

in the event that a first processor element is found to be faulty, removing functionality from the row that contains said first processing element, and allocating functionality to the redundant row.

2. A method as claimed in claim 1, comprising:
allocating the functionality, removed from the row that contains said first processing element, to an adjacent row; and

reallocating the functionality from the adjacent row, to a further row adjacent thereto, as required until functionality has been allocated to the redundant row.

3. A method as claimed in claim 2, wherein the redundant row is located at an edge of the array.

4. A method as claimed in claim 1 or 2, wherein, in operation of said processor array, data is transferred during time slots between processor elements over horizontal buses running between the rows of processor elements and over vertical buses running

between the columns of processor elements, and further
comprising:

when allocating functionality to the processor
elements, such that data is scheduled to be transferred
5   during a first time slot from a first processor element
to a second processor element without using any
vertical bus, reserving said time slot for said data
transfer on a segment of a vertical bus that would be
used in the event of a reallocation of functionality
10   following a determination that either said first
processor element or said second processor element were
faulty.

5.   A method of allocating functionality to a
15   processor array, the processor array comprising a
plurality of processor elements arranged in an array of
rows and columns, the processor elements being
interconnected by buses running between the rows and
columns and by switches located at the intersections of
20   the buses, the method comprising:
in the event that a processor element has a
fault, allocating no functionality to any processor
element in the row. containing said processor element.

25   6.   A processor array, comprising:
a plurality of processor elements arranged in an
array of rows and columns, wherein the arrangement of
processor elements in each row is the same as the
arrangements of processor elements in each other row;
30   pairs of horizontal buses running between the rows
of processor elements, each pair comprising a first
horizontal bus carrying data in a first direction and a
second horizontal bus carrying data in a second
direction opposite to the first direction;

vertical buses running between the columns of processor elements, wherein some pairs of adjacent columns of processor elements have no vertical buses running therebetween, and other pairs of adjacent

5   columns have two buses carrying data in a first direction and two buses carrying data in a second direction opposite to the first direction running therebetween; and

        switches located at the intersections of the

10   horizontal and vertical buses.

       7.   A processor array as claimed in claim 6, wherein each switch comprises:

        a plurality of input buses and a plurality of

15   output buses;

        a memory device, which stores information at each address thereof, indicating what data is to be switched onto each of the output buses; and

        a controller, for counting through addresses of

20   the memory device in a predetermined sequence.

       8.   A processor array as claimed in claim 7, wherein the memory device stores information which indicates whether the data to be switched onto each of

25   the output buses is:

        the data value on one of the input buses;

        the previous data value on said output bus; or

        the value zero.

ABSTRACT

PROCESSOR ARRAY.

An array of processing elements can incorporate a
5   degree of redundancy. Specifically, the array includes
one or more spare, or redundant, rows of array
elements, in addition to the number required to
implement the intended function or functions of the
device. If a defect occurs in one of the processors in
10   the device, then the entire row which includes that
defective processor is not used, and is replaced by a
spare row.

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6
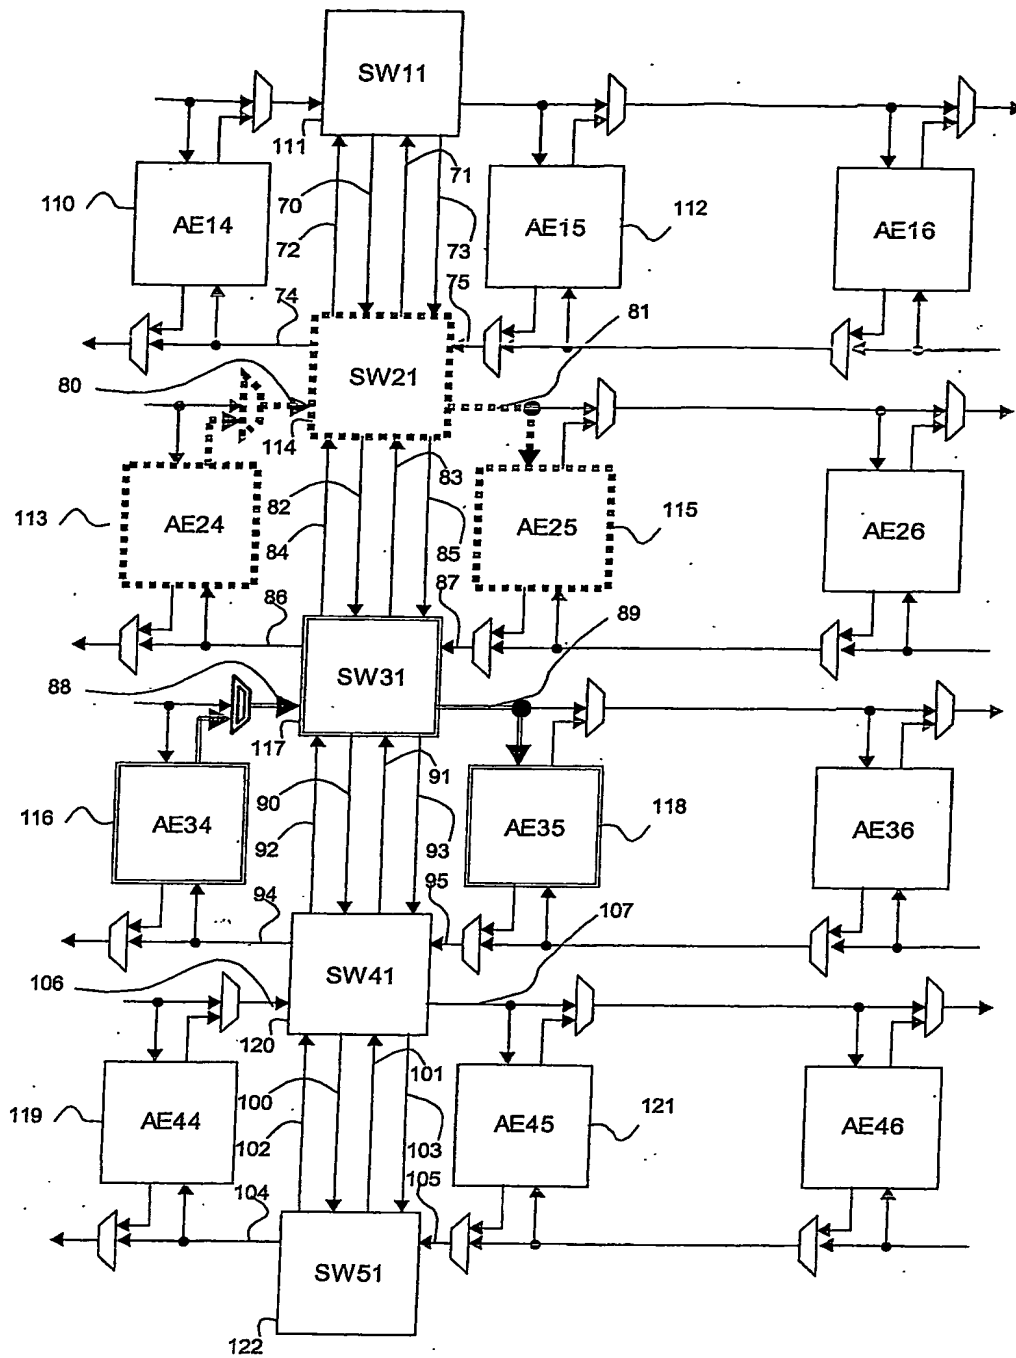
Figure 7
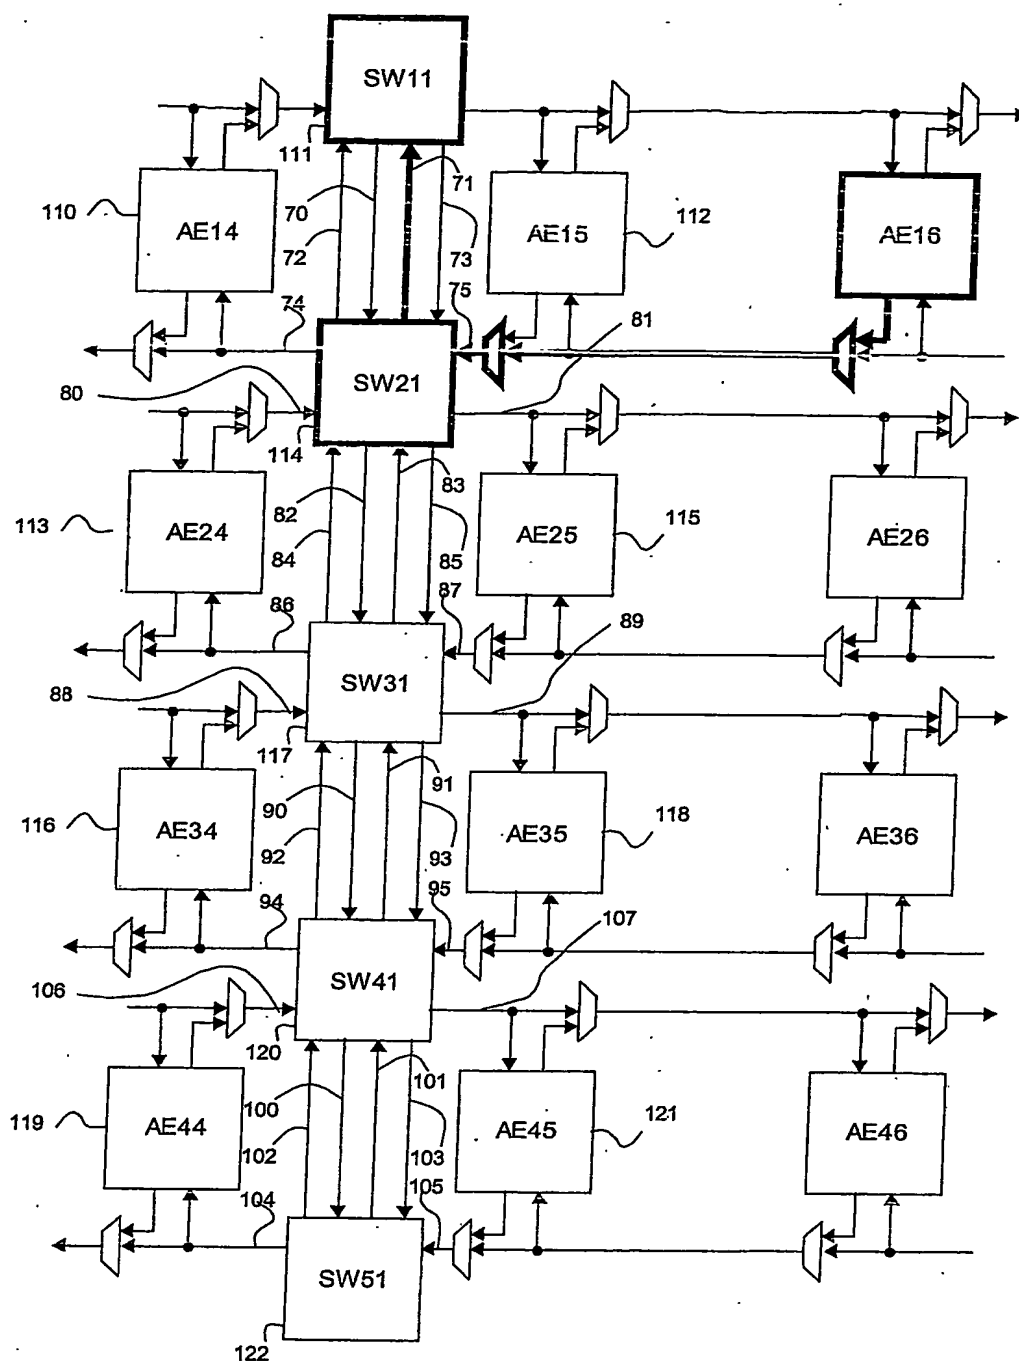
Figure 8
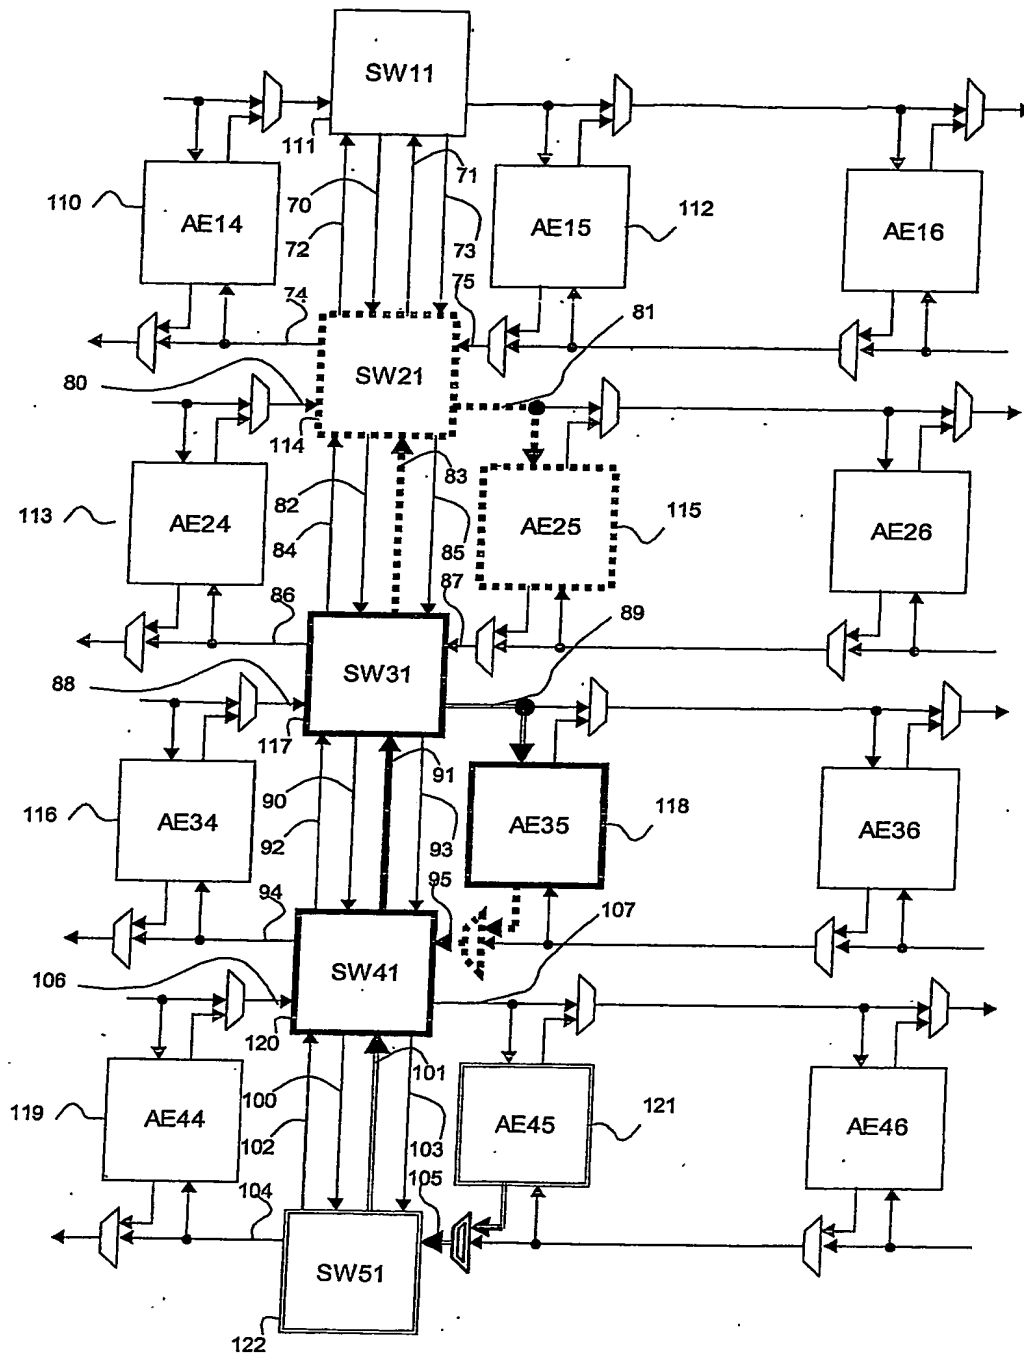
Figure 9
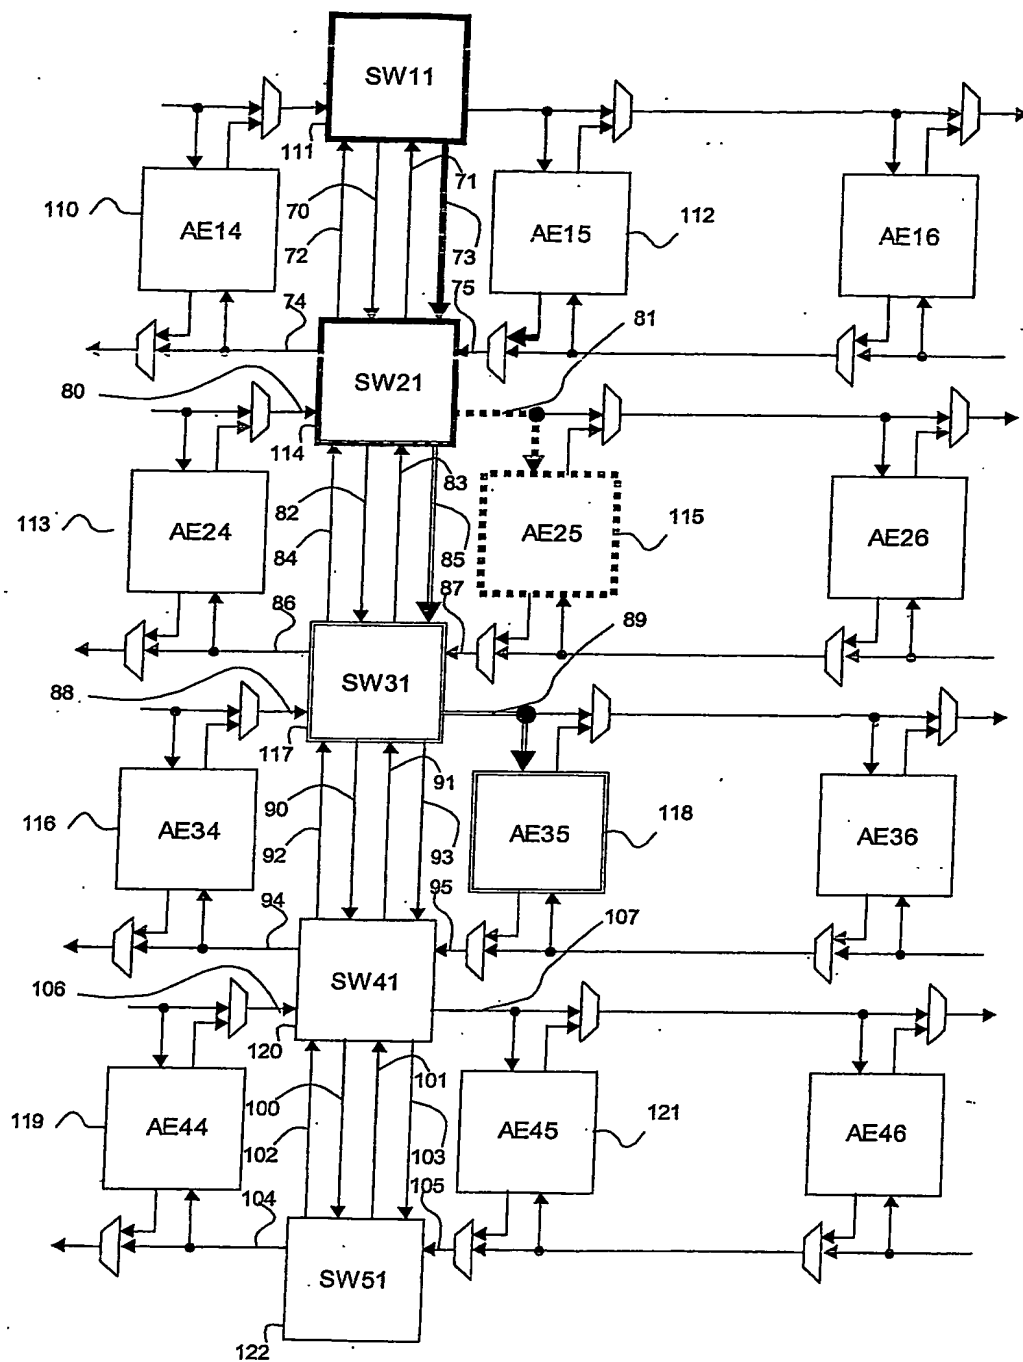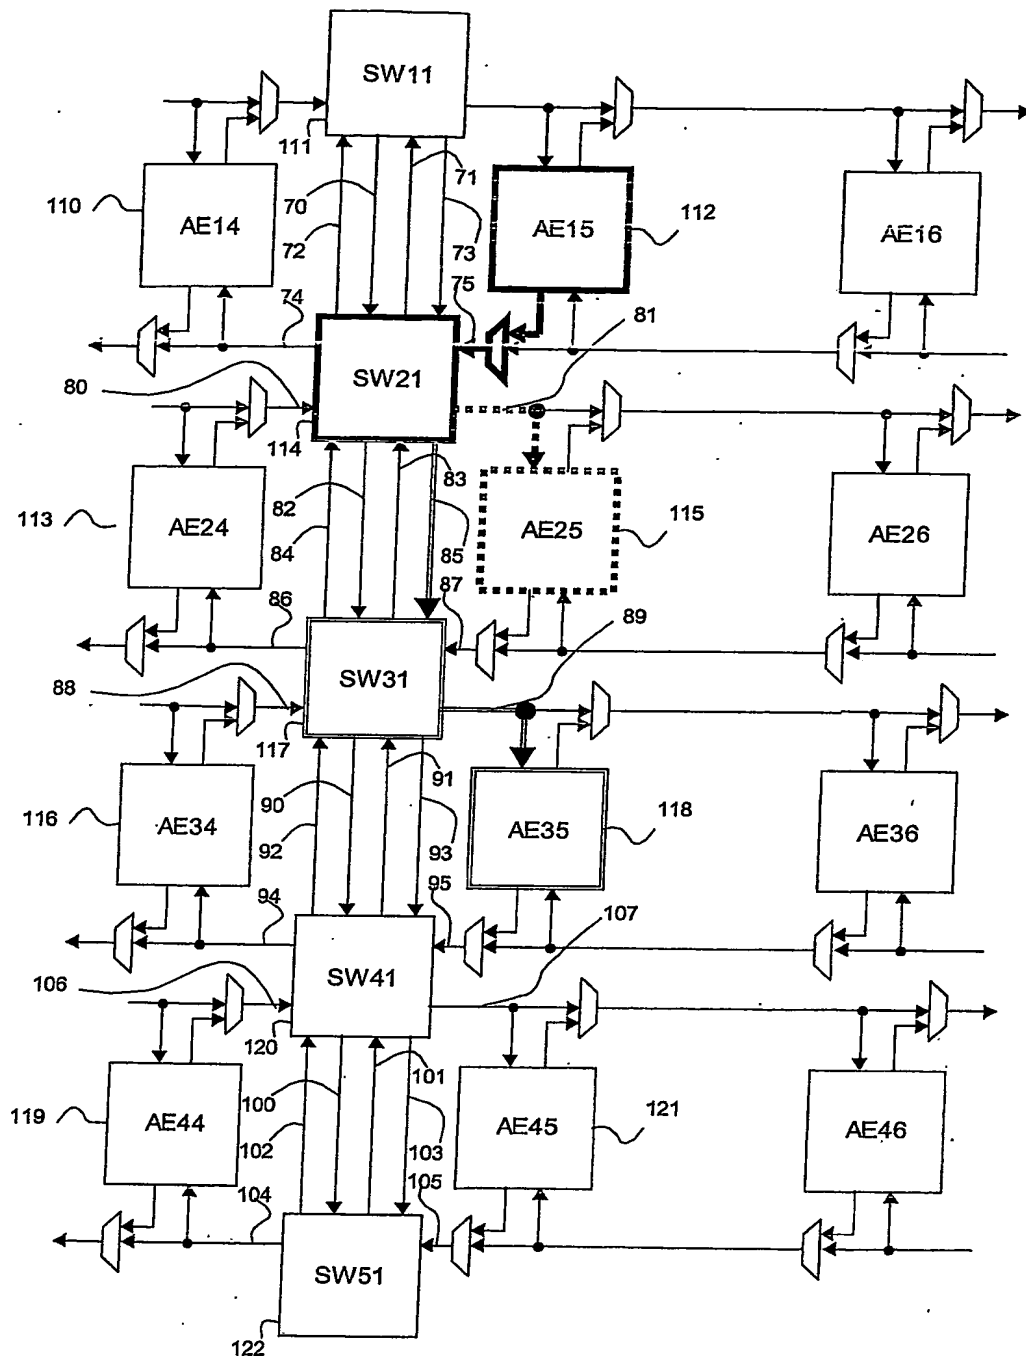
Figure 10

Figure 11
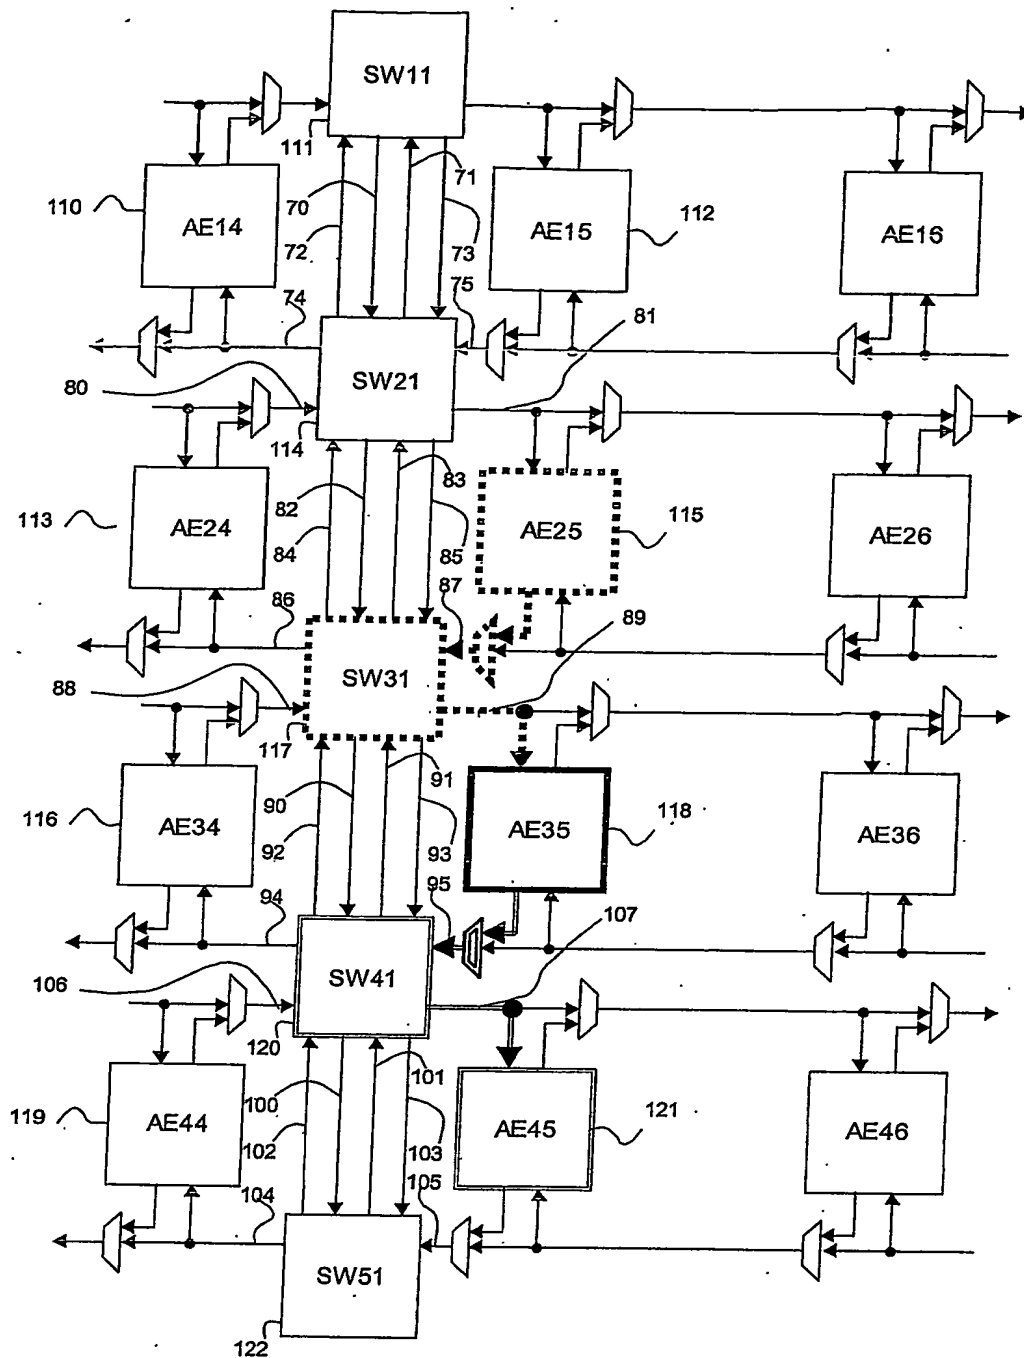
Figure 12
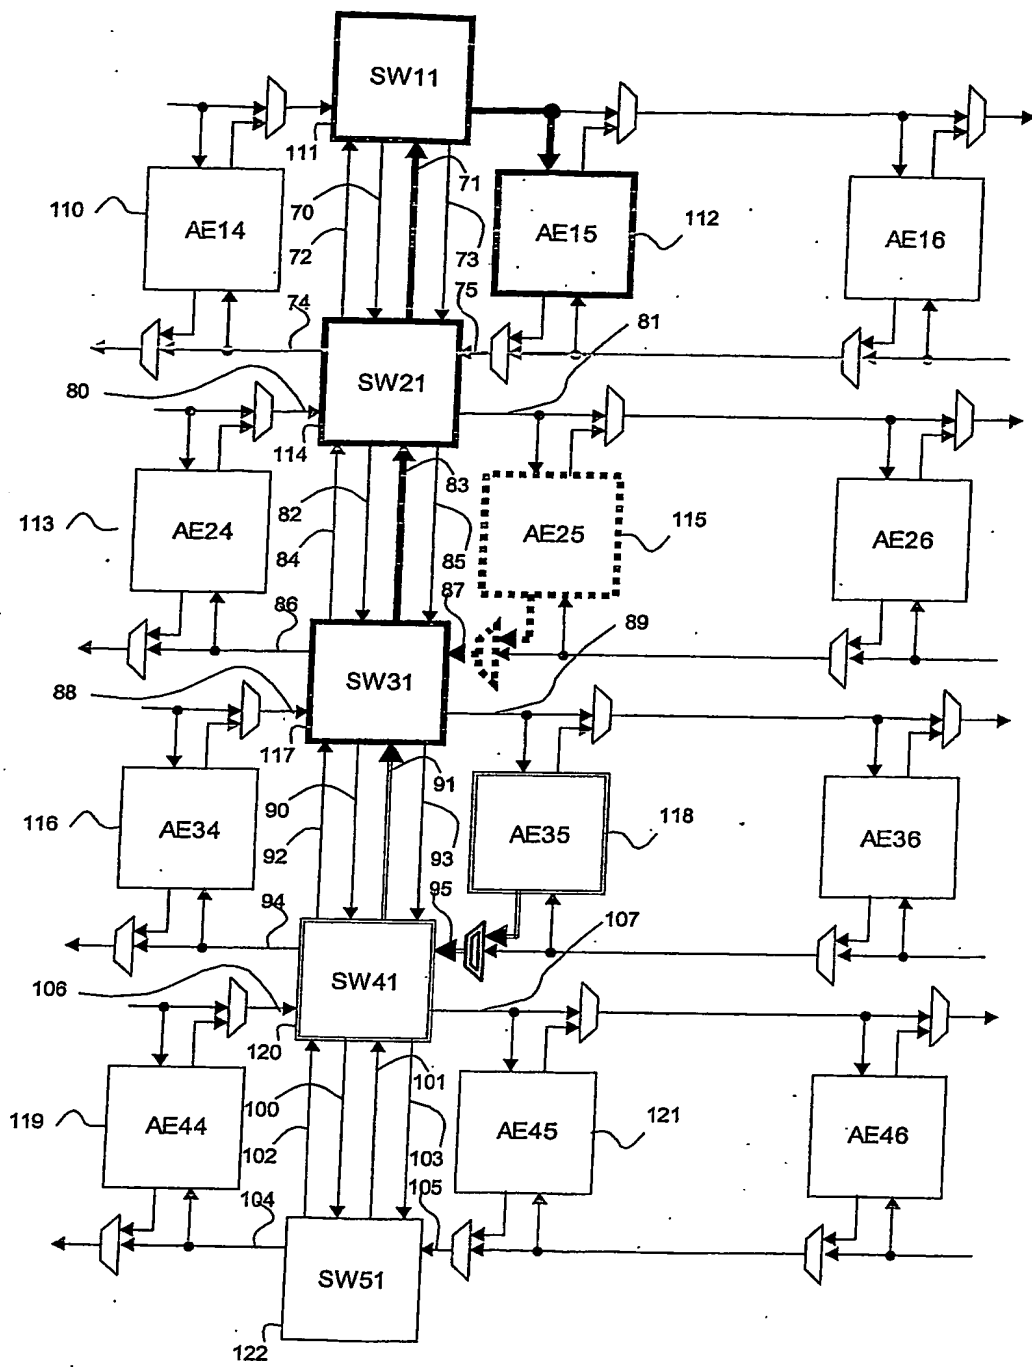
Figure 13
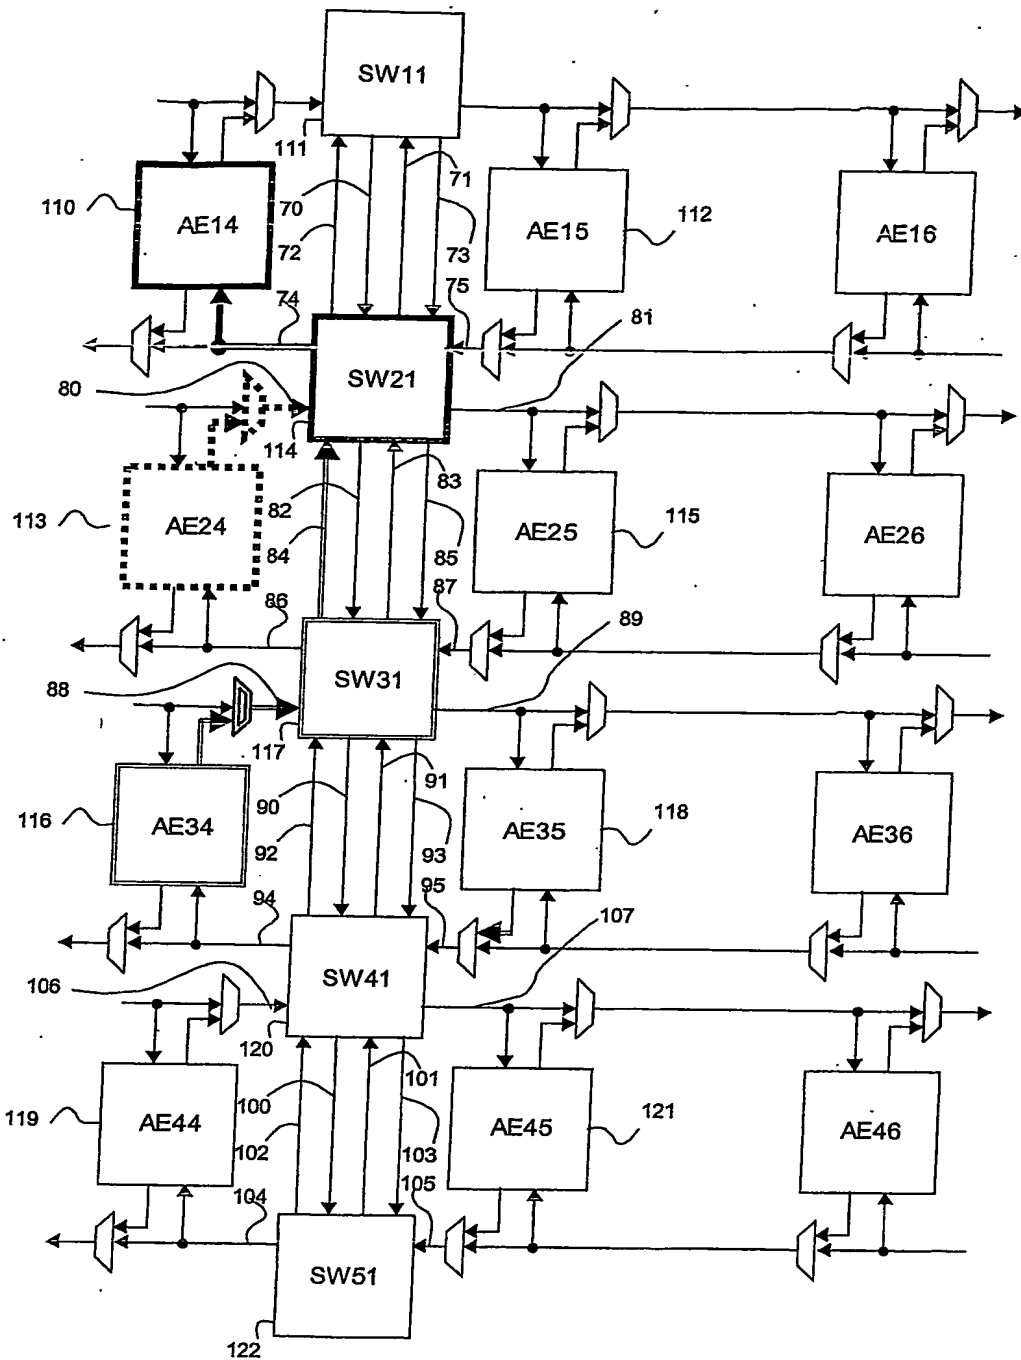
Figure 14
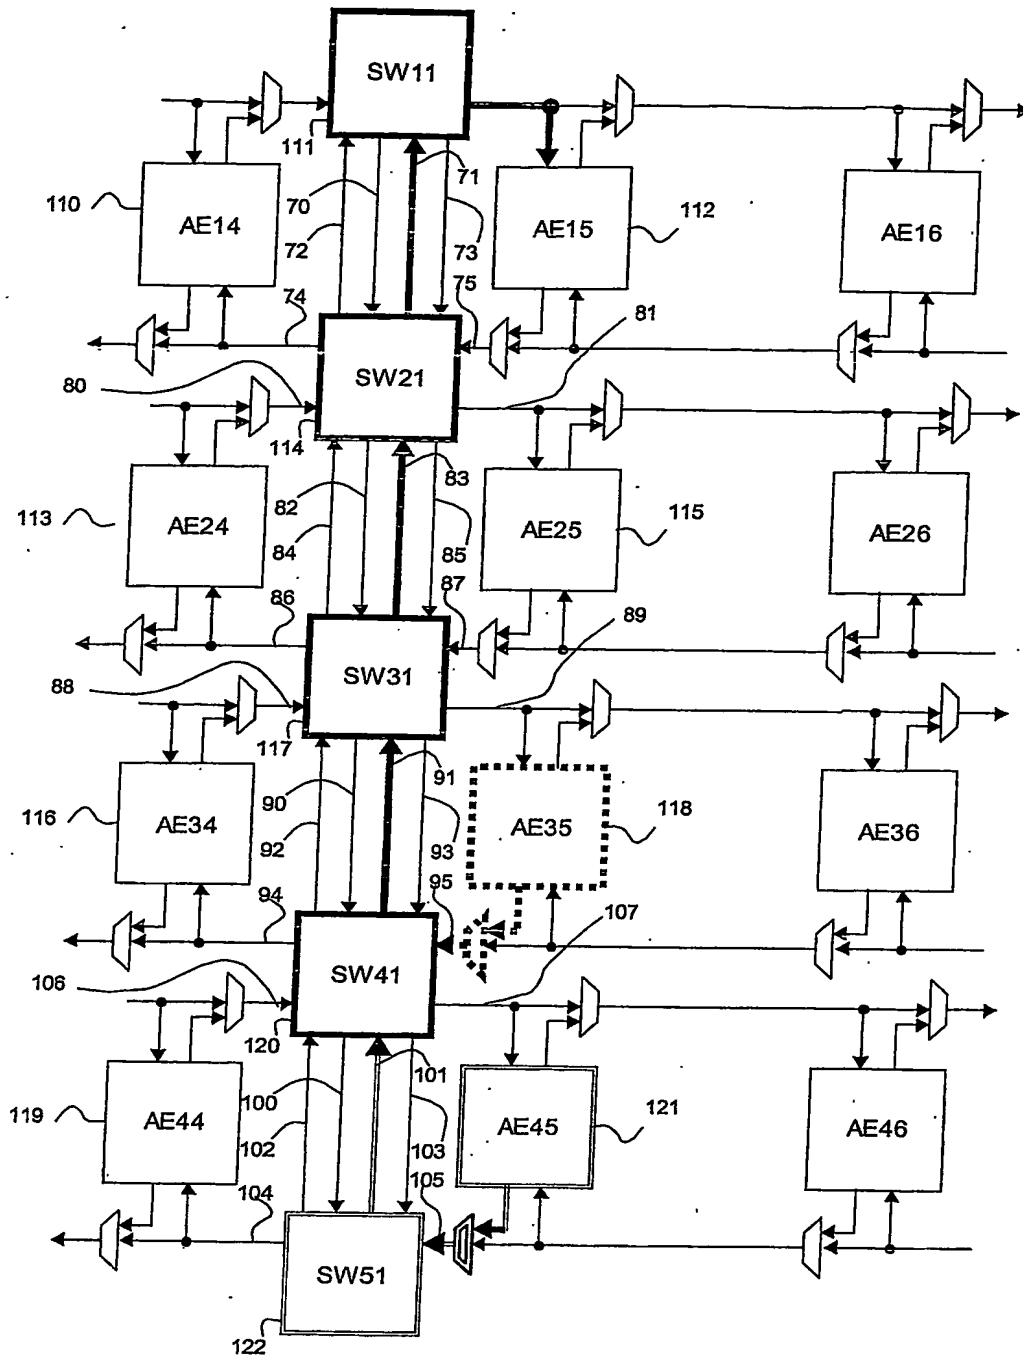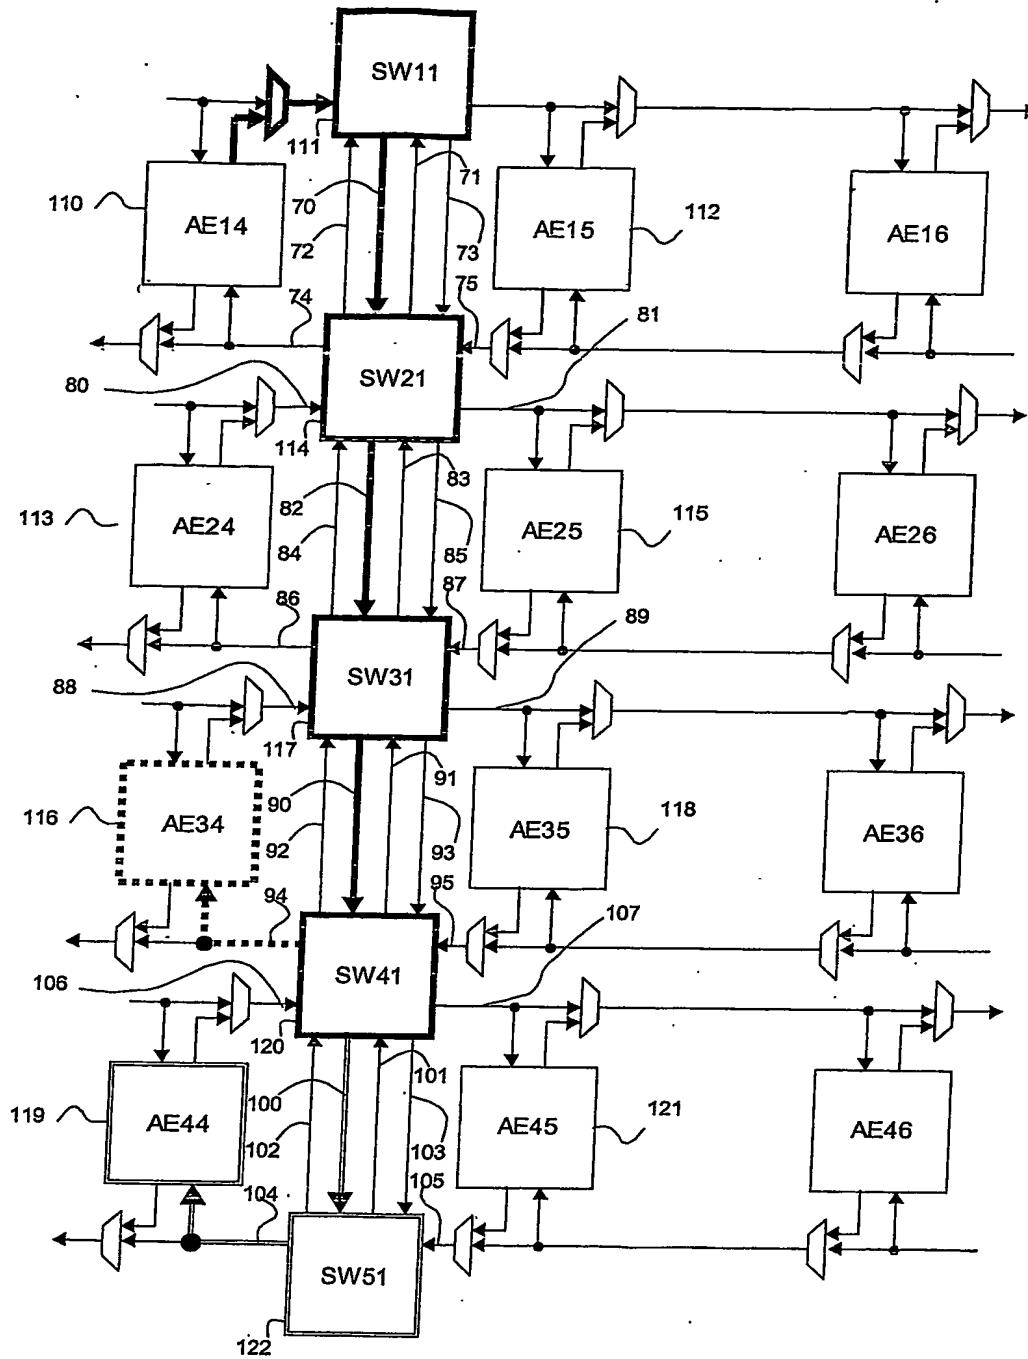
Figure 15

Figure 16

This Page Blank (uspto)